

RANDOMIZATION, SUMS OF SQUARES, AND FASTER REAL ROOT COUNTING FOR TETRANOMIALS AND BEYOND

January 14, 2011

OSBERT BASTANI, CHRISTOPHER J. HILLAR, DIMITAR POPOV, AND J. MAURICE ROJAS

ABSTRACT. Suppose f is a real univariate polynomial of degree D with exactly 4 monomial terms. We present an algorithm, with complexity polynomial in $\log D$ on average (relative to the stable log-uniform measure), for counting the number of real roots of f . The best previous algorithms had complexity super-linear in D . We also discuss connections to sums of squares and \mathcal{A} -discriminants, including explicit obstructions to expressing positive definite sparse polynomials as sums of squares of few sparse polynomials. Our key tool is the introduction of efficiently computable *chamber cones*, bounding regions in coefficient space where the number of real roots of f can be computed easily. Much of our theory extends to n -variate $(n+3)$ -nomials.

1. INTRODUCTION

Counting the real solutions of polynomial equations in one variable is a fundamental ingredient behind many deeper tasks and applications involving the topology of real algebraic sets. However, the intrinsic complexity of this basic enumerative problem becomes a mystery as soon as one considers the input representation in a refined way. Such complexity questions have practical impact for, in many applications such as geometric modelling or the discretization of physically motivated partial differential equations, one encounters polynomials that have sparse expansions relative to some basis. So we focus on new, exponential speed-ups for counting the real roots of certain sparse univariate polynomials of high degree.

Sturm sequences [Stu35], and their later refinements [Hab48, BPR06], have long been a centrally important technique for counting real roots of univariate polynomials. In combination with more advanced algebraic tools such as a Gröbner bases or resultants [GKZ94, BPR06], Sturm sequences have even been applied to algorithmically study the topology of real algebraic sets in arbitrary dimension (see, e.g., [BPR06, Chapters 2, 5, 11, and 16]). However, as we will see below (cf. Examples 1.1 and 1.2), there are obstructions to attaining polynomial intrinsic complexity, for *sparse* polynomials, via Sturm sequences. So we must seek alternatives.

More recently, relating multivariate positive polynomials to sums of squares has become an important algorithmic tool in optimizing real polynomials over semi-algebraic domains [Par03, Las09]. However, there are also obstructions to the use of sums of squares toward speed-ups for sparse polynomials (see Theorem 1.5 below).

Discriminants have a history nearly as long as that of Sturm sequences and sums of squares, but their algorithmic power has not yet been fully exploited. Our main result is that \mathcal{A} -*discriminants* [GKZ94] yield an algorithm for counting real roots, with average-case complexity polynomial in the *logarithm* of the degree, for certain choices of probability distributions on the input (see Theorem 1.4 below). The use of randomization is potentially inevitable in light of the fact that even detecting real roots becomes **NP**-hard already for moderately sparse multivariate polynomials [BRS09, PRT09].

Bastani and Popov were partially supported by NSF REU grant DMS-0552610. Hillar was partially supported by an NSF Postdoctoral Fellowship and an NSA Young Investigator grant. Rojas was partially supported by NSF CAREER Grant DMS-0349309, a Wenner Gren Foundation grant, Sandia National Laboratories, and MSRI.

1.1. From Large Sturm Sequences to Fast Probabilistic Counting. The classical technique of Sturm Sequences [Stu35, BPR06] reduces counting the roots of a polynomial f in a half-open interval $[a, b)$ to a gcd-like computation, followed by sign evaluations for a sequence of polynomials. A key difficulty in these methods, however, is their apparent super-linear dependence on the degree of the underlying polynomial. Consider the following two examples (see also [RY05, Example 1]).

Example 1.1. Setting $f(x_1) := x_1^{317811} - 2x_1^{196418} + 1$, the `realroot` command in Maple 14¹ (which is an implementation of Sturm Sequences) results in an out of memory error after about 31 seconds. The polynomials in the underlying computation, while quite sparse, have coefficients with hundreds of thousands of digits, thus causing this failure. On the other hand, via more recent work [BRS09], one can show that when $c > 0$ and $g(x_1) := x_1^{317811} - cx_1^{196418} + 1$, g has exactly 0, 1, or 2 positive roots according as c is less than, equal to, or greater than $\frac{317811}{(121393^{121393}196418^{196418})^{1/317811}} \approx 1.944\dots$. In particular, our f has exactly 2 positive roots. (We discuss how to efficiently decide the size of monomials in rational numbers with rational exponents in Algorithm 2.15 of Section 2.3 below.) ◇

Example 1.2. Going to tetranomials, consider $f(x_1) := ax_1^{100008} - x_1^{50005} + bx_1^{50004} - 1$ with $a, b > 0$. Then (via the classical Descartes' Rule of Signs [RS02, Cor. 10.1.10, pg. 319]) such an f has exactly 1 or 3 positive roots, but the inequalities characterizing which (a, b) yield either possibility are much more unwieldy than in our last example: there are at least 2, involving polynomials in a and b having tens of thousands of terms. In particular, for $(a, b) = (2, \frac{1}{2})$, Sturm sequences on Maple 14 result in an out of memory error after about 122 seconds. ◇

We have discovered that \mathcal{A} -discriminants, reviewed in Section 2, enable algorithms with complexity polynomial in the *logarithm* of the degree.

Definition 1.3. For any $x = (x_1, \dots, x_d) \in \mathbb{C}^d$, we define $\text{Log}|x| := (\log|x_1|, \dots, \log|x_d|)$ and let the stable log-uniform measure on \mathbb{R}_+^d (resp. \mathbb{Z}^d) be the probability measure ν (resp. ν') defined as follows: $\nu(S) := \lim_{M \rightarrow +\infty} \frac{\mu(\text{Log}|S| \cap [-M, M]^d)}{(2M)^d}$ (resp. $\nu'(S) := \lim_{M \rightarrow +\infty} \frac{\#(\text{Log}|S| \cap \{-M, \dots, M\}^d)}{(2M+1)^d}$), where μ denotes the standard Lebesgue measure on \mathbb{R}^d and $\#(\cdot)$ denotes set cardinality. ◇

Note that the stable log-uniform measure is finitely additive (but not countably additive), and is invariant under reflection across coordinate hyperplanes.

Theorem 1.4. Let $0 < a_2 < a_3 < a_4 = D$ be positive integers,

$$f(x_1) := c_1 + c_2 x_1^{a_2} + c_3 x_1^{a_3} + c_4 x_1^{a_4}$$

with c_1, \dots, c_4 being independent stable log-uniform random variables chosen from \mathbb{R} (resp. \mathbb{Z}), and define $h := \log(2 + \max_i |c_i|)$. Then there is a deterministic algorithm, with complexity polynomial in $\log D$ (resp. h and $\log D$), that computes a number in $\{0, 1, 2, 3\}$ that, with probability 1, is exactly the number of real roots of f . The underlying computational model is the BSS model over \mathbb{R} (resp. the Turing model).

The key idea is that while the regions of coefficients determining polynomials with a constant number of real roots become more complicated as the number of monomial terms increases, one can nevertheless efficiently characterize large subregions — *chamber cones* — where the

¹Running on a 16GB RAM Dell PowerEdge SC1435 departmental server with 2 dual-core Opteron 2212HE 2Ghz processors and OpenSUSE 10.3.

number of real roots is very easy to compute. This motivates the introduction of probability and average-case complexity. The \mathcal{A} -discriminant allows one to make this approach completely precise and algorithmic. In fact, our framework enables us to transparently extend Theorem 1.4 to n -variate $(n+3)$ -nomials (see Theorem 3.18 of Section 3.3).

Our focus on the stable log-uniform measure simplifies our development and has some practical motivation: when one considers N -bit floating-point numbers with uniformly random exponent and mantissa, taking $N \rightarrow +\infty$ and suitably rescaling yields exactly the stable log-uniform measure on \mathbb{Z} . The stable log-uniform measure has also been used in work of Avendaño and Ibrahim to study the expected number of roots of sparse polynomial systems over local fields other than \mathbb{R} [AI11].

It is of course quite natural to ask how the expected complexity in Theorem 1.4 behaves under other well-known measures, e.g., uniform or Gaussian. Unfortunately, the underlying calculations become much more complicated. On a deeper level, it is far from clear what a truly “natural” probability measure on the space of tetranomials is. For instance, for non-sparse polynomials, it is popular to use specially weighted independent Gaussian coefficients since the resulting measure becomes invariant under a natural orthogonal group action (see, e.g., [Kos88, SS96, BSZ00]). However, we are unaware of any study of the types of distributions occurring for the coefficients of polynomials actually occurring in physical applications.

The speed-ups we derive actually hold in far greater generality: see [BRS09, PRT09] for the case of n -variate $(n+k)$ -nomials with $k \leq 2$, Section 3 here for connections to n -variate $(n+3)$ -nomials, the forthcoming paper [AAR11] for the general univariate case, and the forthcoming paper [PRRT11] for chamber cone theory for $n \times n$ sparse polynomial systems. One of the main goals of our paper is thus to illustrate and clarify the underlying theory in a non-trivial special case. We now state our second main result.

1.2. Sparsity and Univariate Sums of Squares. Recent advances in semidefinite programming have produced efficient algorithms for finding sum of squares representations of certain nonnegative polynomials, thus enabling efficient polynomial optimization under certain conditions. When the input is a sparse polynomial it is then natural to ask if there is a sum of squares representation that also respects sparseness. Indeed, it is well-known that a nonnegative univariate polynomial can always be written as a sum of two squares of, usually non-sparse, polynomials (see, e.g., [Pou71] for refinements). The following result demonstrates that a sparse analogue is either unlikely or much more subtle.

Theorem 1.5. *There do not exist absolute constants ℓ and m with the following property: Any trinomial $f \in \mathbb{R}[x_1]$ that is positive on \mathbb{R} can be written in the form $f = g_1^2 + \cdots + g_\ell^2$, for some $g_1, \dots, g_\ell \in \mathbb{R}[x_1]$ with g_i having at most m terms for all i .*

Were there to be a sufficiently efficient representation of positive sparse polynomials as sums of squares, one could then try to use semidefinite programming to find such a representation explicitly for a given polynomial. This in turn could yield an efficient reduction from deciding the existence of real roots to a (small) semidefinite programming problem, similar to the techniques of [Par03]. Our last theorem thus reveals an obstruction to this sums of squares approach.

1.3. Related Approaches. The best known algorithms for real root counting lack speed-ups for sparse polynomials like the average-case complexity bound from our first main result.

For example, in the notation of Theorem 1.4, [LM01] gives an arithmetic complexity bound of $O(D \log^5 D)$ which, via the techniques of [BPR06], yields a bit complexity bound super-linear in $h + D$. No algorithm with complexity polynomial in $\log D$ (deterministic, randomized, or average-case) appears to have been known before for tetranomials. (See [HTZEKM09] for recent speed benchmarks of univariate real solvers.)

As for alternative approaches, softening our concept of sparse sum of squares representation may still enable speed-ups similar to Theorem 1.4 via semidefinite programming. For instance, one could ask if a positive trinomial of degree D always admits a representation as a sum of $\log^{O(1)} D$ squares of polynomials with $\log^{O(1)} D$ terms. This question appears to be completely open.

Example 1.6. *Observe that a quick derivative computation yields that*

$$f(x_1) := x_1^{2^k} - 2^k x_1 + 2^k - 1$$

attains a unique minimum value of 0 at $x = 1$. So this f is nonnegative. On the other hand, one can prove easily by induction that $f(x_1) = 2^{k-1} \sum_{i=0}^{k-1} \frac{1}{2^i} (x_1^{2^i} - 1)^2$, thus yielding an expression for f as a sum of $O(\log D)$ binomials with $D = 2^k$. \diamond

Note also that while we focus on speed-ups that replace the polynomial degree D by $\log D$ in this paper, other practically important speed-ups combining semidefinite programming and sparsity are certainly possible (see, e.g., [Las06, KM09]).

2. BACKGROUND

2.1. Amoebae and Efficient \mathcal{A} -Discriminant Parametrization. Let us first briefly review two important constructions by Gelfand, Kapranov, and Zelevinsky.

Definition 2.1. *Let $c := (c_1, \dots, c_m)$, let $\mathcal{A} = \{a_1, \dots, a_m\} \subset \mathbb{Z}^n$ have cardinality m , and define the corresponding family of (Laurent) polynomials*

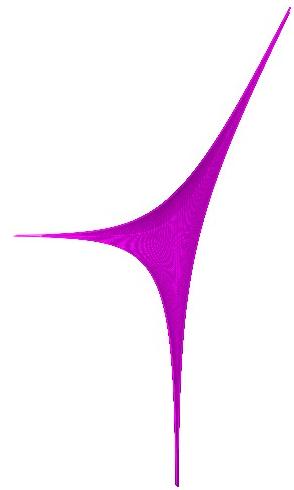
$$\mathcal{F}_{\mathcal{A}} := \{c_1 x^{a_1} + \dots + c_m x^{a_m} \mid c \in \mathbb{C}^m\},$$

where the notation $x^{a_i} := x_1^{a_{1,i}} \cdots x_n^{a_{n,i}}$ is understood. When $c_i \neq 0$ for all $i \in \{1, \dots, m\}$ then we call \mathcal{A} the support of $f(x) = \sum_{i=1}^m c_i x^{a_i}$, also using the notation $\text{Supp}(f)$. \diamond

Definition 2.2. *For any field K we let $K^* := K \setminus \{0\}$. Given any $g \in \mathbb{C}[x_1, \dots, x_n]$, we then define its amoeba, $\text{Amoeba}(g)$, to be $\{\text{Log}|c| \mid c = (c_1, \dots, c_m) \in (\mathbb{C}^*)^m \text{ and } g(c_1, \dots, c_m) = 0\}$. \diamond*

Archimedean Amoeba Theorem. (weaker version of [GKZ94, Cor. 1.8]) *Following the notation of Definition 2.2, the complement of $\text{Amoeba}(g)$ in \mathbb{R}^m is a finite disjoint union of open convex sets. ■*

An example of an amoeba of a bivariate polynomial (see Example 2.5 below) appears in the right-hand illustration. While the complement of the amoeba (in white) appears to have 3 convex connected components, there are in fact 4: the fourth component is a thin sliver emerging further below from the downward pointing tentacle.



Definition 2.3. *Following the notation of Definition 2.1 and letting $f(x) := c_1 x^{a_1} + \dots + c_m x^{a_m}$, we define $\nabla_{\mathcal{A}}$ — the \mathcal{A} -discriminant variety [GKZ94, Chs. 1 & 9–11] — to be the*

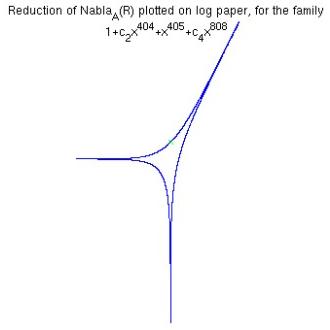
closure of the set of all $[c_1 : \dots : c_m] \in \mathbb{P}_{\mathbb{C}}^{m-1}$ such that

$$f = \frac{\partial f}{\partial x_1} = \dots = \frac{\partial f}{\partial x_n} = 0$$

has a solution in $(\mathbb{C}^*)^n$. We then define (up to sign) the \mathcal{A} -discriminant, $\Delta_{\mathcal{A}} \in \mathbb{Z}[c_1, \dots, c_m]$, to be the (irreducible) defining polynomial of $\nabla_{\mathcal{A}}$ when $\nabla_{\mathcal{A}}$ is a hypersurface. Finally, we let $\nabla_{\mathcal{A}}(\mathbb{R})$ denote the real part of $\nabla_{\mathcal{A}}$. \diamond

Remark 2.4. The $\nabla_{\mathcal{A}}$ considered in this paper will all ultimately be hypersurfaces. \diamond

Example 2.5. Taking $\mathcal{A} = \{0, 404, 405, 808\}$, we see that $\mathcal{F}_{\mathcal{A}}$ consists simply of polynomials of the form $f(x_1) := c_1 + c_2 x_1^{404} + c_3 x_1^{405} + c_4 x_1^{808}$. The underlying \mathcal{A} -discriminant is then a polynomial in the c_i having 609 monomial terms and degree 1604. However, while $\Delta_{\mathcal{A}}$ is unwieldy, we can still easily plot the real part of its zero set $\nabla_{\mathcal{A}}(\mathbb{R})$ via the Horn-Kapranov Uniformization (see its statement below, and the illustration to the right). \diamond



The plotted curve above is the image of the real roots of $\overline{\Delta}_{\mathcal{A}}(c_2, c_4) := \Delta_{\mathcal{A}}(1, c_2, 1, c_4)$ under the $\text{Log}|\cdot|$ map, i.e., the amoeba of $\overline{\Delta}_{\mathcal{A}}$. Amoebae give us a convenient way to introduce polyhedral/tropical methods into our setting. For our last example, the boundary of $\text{Amoeba}(\overline{\Delta}_{\mathcal{A}})$ is contained in the curve above.

\mathcal{A} -discriminants are notoriously large in all but a few restricted settings. For instance, the polynomial $\overline{\Delta}_{\{0,404,405,808\}}$ defining the curve above has the following coefficient for $c_2^{808} c_4$:

903947086576700909448402875044761267196347419431440828445529608410806270
...[2062 digits omitted]... ...93441588472666704061962310429908170311749217550336.

Fortunately, we have the following theorem, describing a one-line parametrization of $\nabla_{\mathcal{A}}$.

The Horn-Kapranov Uniformization. (See [Kap91], [PT05], and [DFS07, Prop. 4.1].) Given $\mathcal{A} := \{a_1, \dots, a_m\} \subset \mathbb{Z}^n$ with $\nabla_{\mathcal{A}}$ a hypersurface, the discriminant locus $\nabla_{\mathcal{A}}$ is exactly the closure of

$$\{[u_1 \lambda^{a_1} : \dots : u_m \lambda^{a_m}] \mid u := (u_1, \dots, u_m) \in \mathbb{C}^m, \mathcal{A}u = \mathbf{0}, \sum_{i=1}^m u_i = 0, \lambda = (\lambda_1, \dots, \lambda_n) \in (\mathbb{C}^*)^n\}. \blacksquare$$

Thus, once we know the null-space of an $(n+1) \times m$ matrix, we have a formula parametrizing $\nabla_{\mathcal{A}}$. Recall that for any two subsets $U, V \subseteq \mathbb{R}^N$, their Minkowski sum $U + V$ is $\{u + v \mid u \in U, v \in V\}$. Also, for any matrix M , we let M^T denote its transpose.

Corollary 2.6. Following the notation above, let $\hat{\mathcal{A}}$ denote the $(n+1) \times m$ matrix whose i^{th} column has coordinates corresponding to $1 \times a_i$, let $B \in \mathbb{R}^{m \times p}$ be any real matrix whose columns are a basis for the right null-space of $\hat{\mathcal{A}}$, and define $\varphi : \mathbb{C}^p \rightarrow \mathbb{R}^m$ via $\varphi(t) := \log |tB^T|$. Then $\text{Amoeba}(\Delta_{\mathcal{A}})$ is the Minkowski sum of the row space of $\hat{\mathcal{A}}$ and $\varphi(\mathbb{C}^p)$. \blacksquare

If one is familiar with elimination theory, then it is evident from the Horn-Kapranov Uniformization that discriminant amoebae are subspace bundles over a lower-dimensional amoeba. This is a geometric reformulation of the homogeneities satisfied by the polynomial $\Delta_{\mathcal{A}}$.

Example 2.7. Continuing Example 2.5, we observe that $\hat{\mathcal{A}} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 404 & 405 & 808 \end{bmatrix}$ has right null-space generated by the respective transposes of $(1, -405, 404, 0)$ and $(1, -2, 0, 1)$. The Horn-Kapranov Uniformization then tells us that $\nabla_{\mathcal{A}}$ is simply the closure of the rational

surface $\{[t_1 + t_2 : (-405t_1 - 2t_2)\lambda^{404} : 404t_1\lambda^{405} : t_2\lambda^{808}] \mid t_1, t_2 \in \mathbb{C}, \lambda \in \mathbb{C}^*\} \subset \mathbb{P}_{\mathbb{C}}^3$. Note that f and $\frac{1}{c_1}f$ have the same roots and that $u \mapsto u^{1/405}$ is a well-defined bijection on \mathbb{R} that preserves sign. Note also that the roots of f and $\bar{f}(y) := \frac{1}{c_1}f\left(\left(\frac{c_1}{c_3}\right)^{1/405}y\right)$ differ only by a scaling when f has real coefficients, and that \bar{f} is of the form $1 + c'_2y^{404} + y^{405} + c'_4y^{808}$. It then becomes clear that we can reduce the study of $\nabla_{\mathcal{A}}(\mathbb{R})$ to a lower-dimensional slice: intersecting $\nabla_{\mathcal{A}}$ with the plane defined by $c_1 = c_3 = 1$ yields the following parametrized curve in \mathbb{C}^2 :

$$\overline{\nabla}_{\mathcal{A}} := \left\{ \left(\frac{-405t_1 - 2t_2}{t_1 + t_2} \left(\frac{404t_1}{t_1 + t_2} \right)^{-404/405}, \frac{t_2}{t_1 + t_2} \left(\frac{404t_1}{t_1 + t_2} \right)^{-808/405} \right) \mid t_1, t_2 \in \mathbb{C} \right\}.$$

In other words, the preceding curve is the closure of the set of all $(c'_2, c'_4) \in (\mathbb{C}^*)^2$ such that $1 + c'_2x^{404} + x^{405} + c'_4x^{808}$ has a degenerate root in \mathbb{C}^* . Our preceding illustration of the image of $\overline{\nabla}_{\mathcal{A}}(\mathbb{R})$ within $\text{Amoeba}(\overline{\Delta}_{\mathcal{A}})$ (after taking log absolute values of coordinates) thus has the following explicit parametrization:

$$\left\{ \left(\log |405t_1 + 2t_2| - \frac{1}{405} \log |t_1 + t_2| - \frac{404}{405} \log |404t_1|, \log |t_2| + \frac{403}{405} \log |t_1 + t_2| - \frac{808}{405} \log |404t_1| \right) \right\}_{[t_1:t_2] \in \mathbb{P}_{\mathbb{R}}^1} \diamond$$

A geometric fact about amoebae that will prove quite useful here is the following elegant quantitative result of Passare and Rullgård. Recall that the *Newton polytope* of a Laurent polynomial $f \in \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ is the *convex hull* of² the exponent vectors appearing in the monomial term expansion of f .

Passare-Rullgård Theorem. [PR04, Cor. 1] Suppose $f \in \mathbb{C}[x_1^{\pm 1}, x_2^{\pm 1}]$ has Newton polygon P . Then $\text{Area}(\text{Amoeba}(f)) \leq \pi^2 \text{Area}(P)$. ■

2.2. Discriminant Chambers and Cones. \mathcal{A} -discriminants are central in real root counting because the real part of $\nabla_{\mathcal{A}}$ determines where, in coefficient space, the real zero set of a polynomial changes topology. Recall that a *cone* in \mathbb{R}^m is any subset closed under addition and nonnegative linear combinations. The dimension of a cone C is the dimension of the smallest flat containing C .

Definition 2.8. Suppose $\mathcal{A} = \{a_1, \dots, a_m\} \subset \mathbb{Z}^n$ has cardinality m and $\nabla_{\mathcal{A}}$ is a hypersurface. We then call any connected component \mathcal{C} of the complement of $\nabla_{\mathcal{A}}$ in $\mathbb{P}_{\mathbb{R}}^{m-1} \setminus \{c_1 \cdots c_m = 0\}$ a (real) discriminant chamber. Also let \hat{A} denote the $(n+1) \times m$ matrix whose i^{th} column has coordinates corresponding to $1 \times a_i$ and let B be any real matrix $B = [b_{i,j}] \in \mathbb{R}^{m \times p}$ with $\begin{bmatrix} \hat{A} \\ B^T \end{bmatrix}$ invertible. If $\log |\mathcal{C}|B$ contains an m -dimensional cone then we call \mathcal{C} an outer chamber (of $\nabla_{\mathcal{A}}$). All other chambers of $\nabla_{\mathcal{A}}$ are called inner chambers (of $\nabla_{\mathcal{A}}$). Finally, we call the formal expression $(c_1, \dots, c_m)^B := \left(c_1^{b_{1,1}} \cdots c_m^{b_{m,1}}, \dots, c_1^{b_{1,p}} \cdots c_m^{b_{m,p}} \right)$ a monomial change of variables, and we refer to images of the form \mathcal{C}^B (with \mathcal{C} an inner or outer chamber) as reduced chambers. ◇

It is easily verified that $\log |\mathcal{C}^B| = \log |\mathcal{C}|B$. The latter notation simply means the image of $\log |\mathcal{C}|$ under right multiplication by the matrix B .

Example 2.9. The illustration from Example 2.5 shows \mathbb{R}^2 partitioned into what appear to be 3 convex and unbounded regions, and 1 non-convex unbounded region. There are in fact 4 convex and unbounded regions, with the fourth visible only if the downward pointed spike were allowed to extend much farther down. So $\mathcal{A} = \{0, 404, 405, 808\}$ results in exactly 4 reduced outer chambers. ◇

²i.e., smallest convex set containing...

Note that exponentiating by any B as above yields a well-defined multiplicative homomorphism from $(\mathbb{R}^*)^m$ to $(\mathbb{R}^*)^p$ when B has rational entries with all denominators odd. In particular, the definition of outer chamber is independent of B , since (for the B considered above) $\text{Log}|\mathcal{C}^B|$ is unbounded and convex iff $\text{Log}|\mathcal{C}^{B^*}|$ is unbounded and convex, where B^* is any matrix whose columns are a basis for the orthogonal complement of the row space of $\hat{\mathcal{A}}$.

One can in fact reduce the study of the topology of sparse polynomial real zero sets to representatives coming from reduced discriminant chambers. A special case of this reduction is the following result.

Lemma 2.10. (See [DRRS07, Prop. 2.17].) *Suppose $\mathcal{A} \subset \mathbb{Z}^n$ has $n+3$ elements, \mathcal{A} is not contained in any $(n-1)$ -flat, $\mathcal{A} \cap Q$ has cardinality n for all facets Q of $\text{Conv}\mathcal{A}$, all the entries of $B \in \mathbb{Q}^{(n+3) \times 2}$ have odd denominator, and $\begin{bmatrix} \hat{\mathcal{A}} \\ B^T \end{bmatrix}$ is invertible. Also let $f, g \in \mathcal{F}_{\mathcal{A}} \setminus \nabla_{\mathcal{A}}$ have respective real coefficient vectors c and c' with c^B and c'^B lying in the same reduced discriminant chamber. Then all the complex roots of f and g are non-singular, and the respective zero sets of f and g in $(\mathbb{R}^*)^n$ are diffeotopic. In particular, for $n=1$, we have that either $f(x)$ and $g(x)$ have the same number of positive roots or $f(x)$ and $g(-x)$ have the same number of positive roots. ■*

2.3. Integer Linear Algebra and Linear Forms in Logarithms. Certain quantitative results on integer matrix factorizations and linear forms in logarithms will prove crucial for our main algorithmic results.

Definition 2.11. Let $\mathbb{Z}^{n \times m}$ denote the set of $n \times m$ matrices with all entries integral, and let $\mathbf{GL}_n(\mathbb{Z})$ denote the set of all matrices in $\mathbb{Z}^{n \times n}$ with determinant ± 1 (the set of unimodular matrices). Recall that any $n \times m$ matrix $[u_{ij}]$ with $u_{ij} = 0$ for all $i > j$ is called upper triangular.

Given any $M \in \mathbb{Z}^{n \times m}$, we then call an identity of the form $UM = H$, with $H = [h_{ij}] \in \mathbb{Z}^{n \times m}$ upper triangular and $U \in \mathbf{GL}_n(\mathbb{Z})$, a Hermite factorization of M . Also, if we have the following conditions in addition:

- (1) $h_{ij} \geq 0$ for all i, j .
- (2) for all i , if j is the smallest j' such that $h_{ij'} \neq 0$ then $h_{ij} > h_{i'j}$ for all $i' \leq i$.

then we call H the Hermite normal form of M . ◇

Proposition 2.12. We have that $x^{AB} = (x^A)^B$ for any $A, B \in \mathbb{Z}^{n \times n}$. Also, for any field K , the map defined by $m(x) = x^U$, for any unimodular matrix $U \in \mathbb{Z}^{n \times n}$, is an automorphism of $(K^*)^n$. ■

Theorem 2.13. [Sto00, Ch. 6, Table 6.2, pg. 94] For any $A = [a_{i,j}] \in \mathbb{Z}^{n \times m}$ with $m \geq n$, the Hermite factorization of A can be computed within $O(nm^{2.376} \log^2(m \max_{i,j} |a_{i,j}|))$ bit operations. Furthermore, the entries of all matrices in the Hermite factorization have bit size $O(m \log(m \max_{i,j} |a_{i,j}|))$. ■

The following result is a very special case of work of Nesterenko that dramatically refines Baker's famous theorem on linear forms in logarithms [Bak77].

Theorem 2.14. (See [Nes03, Thm. 2.1, Pg. 55].) For any integers $\gamma_1, \alpha_1, \dots, \gamma_N, \alpha_N$ with $\alpha_i \geq 2$ for all i , define $\Lambda(\gamma, \alpha) := \gamma_1 \log(\alpha_1) + \dots + \gamma_N \log(\alpha_N)$. Then $\Lambda(\gamma, \alpha) \neq 0 \implies \log \left| \frac{1}{\Lambda(\gamma, \alpha)} \right|$ is bounded above by $2.9(N+2)^{9/2}(2e)^{2N+6}(2 + \log \max_j |\gamma_j|) \prod_{j=1}^N \log \alpha_j$. ■

The most obvious consequence of lower bounds for linear forms in logarithms is an efficient way to determine the signs of monomials in integers.

Algorithm 2.15.

Input: Positive integers $\alpha_1, u_1, \dots, \alpha_M, u_M$ and $\beta_1, v_1, \dots, \beta_N, v_N$ with $\alpha_i, \beta_i \geq 2$ for all i .

Output: The sign of $\alpha_1^{u_1} \cdots \alpha_M^{u_M} - \beta_1^{v_1} \cdots \beta_N^{v_N}$.

Description:

- (0) Check via gcd-free bases (see, e.g., [BS96, Sec. 8.4]) whether $\alpha_1^{u_1} \cdots \alpha_M^{u_M} = \beta_1^{v_1} \cdots \beta_N^{v_N}$. If so, output “They are equal.” and STOP.

- (1) Let $U := \max\{u_1, \dots, u_M, v_1, \dots, v_N\}$ and

$$E := \frac{2.9}{\log 2} (2e)^{2M+2N+6} (1 + \log U) \times \left(\prod_{i=1}^M \log \alpha_i \right) \left(\prod_{i=1}^N \log \beta_i \right).$$

- (2) For all $i \in [M]$ (resp. $i \in [N]$), let A_i (resp. B_i) be a rational number agreeing with $\log \alpha_i$ (resp. $\log \beta_i$) in its first $2 + E + \log_2 M$ (resp. $2 + E + \log_2 N$) leading bits.³

- (3) Output the sign of $\left(\sum_{i=1}^M u_i A_i \right) - \left(\sum_{i=1}^N v_i B_i \right)$ and STOP.

Lemma 2.16. Algorithm 2.15 is correct and, following the preceding notation, runs within a number of bit operations asymptotically linear in

$$(M+N)(30)^{M+N} L(\log U) \left(\prod_{i=1}^M L(\log \alpha_i) \right) \left(\prod_{i=1}^N L(\log \beta_i) \right),$$

where $L(x) := x(\log x)^2 \log \log x$. ■

Lemma 2.16 follows immediately from Theorem 2.14, the well-known fast iterations for approximating \log (see, e.g., [Ber03]), and the known refined bit complexity estimates for fast multiplication (see, e.g., [BS96, Table 3.1, pg. 43]).

3. CHAMBER CONES AND POLYHEDRAL MODELS

3.1. Defining and Describing Chamber Cones.

Definition 3.1. Suppose $X \subset \mathbb{R}^m$ is convex and $Q \supseteq X$ is the polyhedral cone consisting of all $c \in \mathbb{R}^m$ with $c + X \subseteq X$. We call Q the recession cone for X and, if $p \in \mathbb{R}^m$ satisfies (1) $p + Q \supseteq X$ and (2) $p + c + Q \not\supseteq X$ for any $c \in Q \setminus \{\mathbf{0}\}$, then we call $p + Q$ the placed recession cone. In particular, the placed recession cone for $\text{Log}|\mathcal{C}|$ for \mathcal{C} an outer chamber (resp. a reduced outer chamber) is called a chamber cone (resp. a reduced chamber cone) of $\nabla_{\mathcal{A}}$. We call the facets of the (reduced) chamber cones of $\nabla_{\mathcal{A}}$ (reduced) walls of $\nabla_{\mathcal{A}}$. We also refer to walls of dimension 1 as rays. ◇

³For definiteness, let us use Arithmetic-Geometric Mean Iteration as in [Ber03] to find these approximations.

Example 3.2. Returning to Example 2.5, let us draw the rays that are the boundaries of the 4 reduced chamber cones: While there appear to be just 3 reduced chamber cones, there are in fact 4, since there is an additional (slender) reduced chamber cone with vertex placed much farther down. (The magnified illustration to the right actually shows 2 nearly parallel rays going downward, very close together.) Note also that reduced chamber cones need not share vertices. \diamond

Chamber cones are well-defined since chambers are *log-convex*, being the domains of convergence of a particular class of hypergeometric series (see, e.g., [GKZ94, Ch. 6]). A useful corollary of the Horn-Kapranov Uniformization is a surprisingly simple and explicit description for chamber cones.

Definition 3.3. Suppose $\mathcal{A} \subset \mathbb{Z}^n$ has cardinality $n+3$, is not contained in any $(n-1)$ -flat, and is not a pyramid.⁴ Also let B be any real $(n+3) \times 2$ matrix whose columns are a basis for the right null space of $\hat{\mathcal{A}}$ and let $\beta_1, \dots, \beta_{n+3}$ be the rows of B . We then call any set of indices $\mathcal{I} \subset \{1, \dots, n+3\}$ satisfying:

- (a) $[\beta_i]_{i \in \mathcal{I}}$ is a maximal rank 1 submatrix of B .
- (b) $\sum_{i \in \mathcal{I}} \beta_i$ is not the zero vector.

a radiant subset corresponding to \mathcal{A} . \diamond

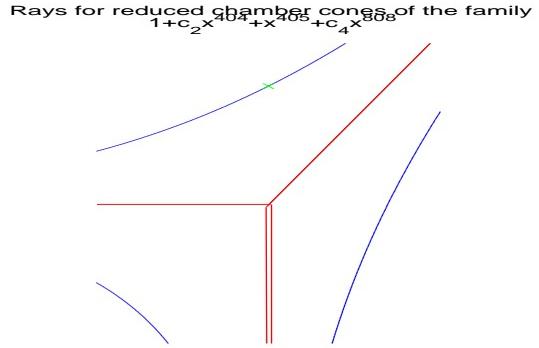
Theorem 3.4. Suppose $\mathcal{A} \subset \mathbb{Z}^n$ has cardinality $n+3$, is not contained in any $(n-1)$ -flat, is not a pyramid, and $\nabla_{\mathcal{A}}$ is a hypersurface. Also let B be any real $(n+3) \times 2$ matrix whose columns are a basis for the right null space of $\hat{\mathcal{A}}$ and let $\beta_1, \dots, \beta_{n+3}$ be the rows of B .

Finally, let $S := [s_{i,j}] := B \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} B^T$ and let s_i denote the row vector whose j^{th} coordinate is 0 or $\log |s_{i,j}|$ according as $s_{i,j}$ is 0 or not. Then each wall of $\nabla_{\mathcal{A}}$ is the Minkowski sum of the row-space of $\hat{\mathcal{A}}$ and a ray of the form $s_i - \mathbb{R}_+ \sum_{j \in \mathcal{I}} e_j$ for some unique radiant subset \mathcal{I} of \mathcal{A} and any $i \in \mathcal{I}$. In particular, the number of walls of $\nabla_{\mathcal{A}}$, the number of chamber cones of $\nabla_{\mathcal{A}}$, and the number of radiant subsets corresponding to \mathcal{A} are all identical, and at most $n+3$.

Note that the hypotheses on \mathcal{A} are trivially satisfied when $\mathcal{A} \subset \mathbb{Z}$ has cardinality 4. Note also that the definition of a radiant subset corresponding to \mathcal{A} is independent of the chosen basis B since the definition is invariant under column operations on B .

Remark 3.5. Theorem 3.4 thus refines an earlier result of Dickenstein, Feichtner, and Sturmfels ([DFS07, Thm. 1.2]) where, in essence, unshifted variants of chamber cones (all going through the origin) were computed for nonpyramidal $\mathcal{A} \subset \mathbb{Z}^n$ of arbitrary cardinality. A version of Theorem 3.4 for \mathcal{A} of arbitrary cardinality will appear in [PRRT11]. \diamond

Example 3.6. It is easy to show that a generic \mathcal{A} satisfying the hypotheses of our theorem will have exactly $n+3$ chamber cones, e.g., Example 2.5. It is also almost as easy to construct examples having fewer chamber cones. For instance, taking $\mathcal{A} := \left\{ \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \end{bmatrix} \right\}$,



⁴The last assumption simply means that there is no point $a \in \mathcal{A}$ such that $\mathcal{A} \setminus \{a\}$ lies in an $(n-1)$ -flat.

we see that $B := \begin{bmatrix} -1 & 0 \\ 2 & 1 \\ -2 & -2 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$ satisfies the hypotheses of Theorem 3.4 and that $\{1, 5\}$ is a non-radiant subset. So the underlying $\nabla_{\mathcal{A}}$ has only 3 chamber cones. \diamond

Proof of Theorem 3.4: First note that by Corollary 2.6, $\text{Amoeba}(\Delta_{\mathcal{A}})$ is the Minkowski sum of $\varphi(\mathbb{C}^2)$ and the row space of $\hat{\mathcal{A}}$, where $\varphi(t) := \text{Log}|tB^T|$. So then, determining the walls reduces to determining the directions orthogonal to the row space of $\hat{\mathcal{A}}$ in which $\varphi(t)$ becomes unbounded.

Since $\mathbf{1} := (1, \dots, 1)$ is in the row space of $\hat{\mathcal{A}}$, we have $\mathbf{1}B = \mathbf{0}$ and thus $\varphi(t) = \varphi(t/M)$ for all $M > 0$. So we can restrict to the compact subset $\{(t_1, t_2) \mid |t_1|^2 + |t_2|^2 = 1\}$ and observe that $\varphi(t)$ becomes unbounded iff $t\beta_i^T \rightarrow 0$ for some i . In particular, we see that there are no more than $n+3$ reduced walls. Note also that $t\beta_i^T \rightarrow 0$ iff t tends to a suitable (nonzero) multiple of $\beta_i \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, in which case the coordinates of $\varphi(t)$ becoming unbounded are exactly those with index $j \in \mathcal{I}$ where \mathcal{I} is the unique radiant subset corresponding to those rows of B that are nonzero multiples of b_i . (The assumption that \mathcal{A} not be a pyramid implies that B can have no zero rows.) Furthermore, the coordinates of $\varphi(t)$ that become unbounded each tend to $-\infty$. Note that Condition (b) of the radiance condition comes into play since we are looking for directions *orthogonal to the row-space of $\hat{\mathcal{A}}$* in which $\varphi(t)$ becomes unbounded.

We thus obtain that each wall is of the asserted form. However, we still need to account for the coordinates of $\varphi(t)$ that remain bounded. Now, if t tends to a suitable (nonzero) multiple of $\beta_i \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$, then it clear that any coordinate of $\varphi(t)$ of index $j \notin \mathcal{I}$ tends to $s_{i,j}$ (modulo a multiple of $\mathbf{1}$ added to $\varphi(t)$). So we have indeed described every wall, and given a bijection between radiant subsets corresponding to \mathcal{A} and the walls of $\nabla_{\mathcal{A}}$.

To conclude, note that the row space of $\hat{\mathcal{A}}$ has dimension $n+1$ by construction, so the walls are all actually (parallel) n -plane bundles over rays. So by the Archimedean Amoeba Theorem, each outer chamber of $\nabla_{\mathcal{A}}$ must be bounded by 2 walls and the walls have a natural cyclic ordering. Thus the number of chamber cones is the same as the number of rays and we are done. ■

3.2. Which Chamber Cone Contains Your Problem? An important consequence of Theorem 3.4 is that while the underlying \mathcal{A} -discriminant polynomial $\Delta_{\mathcal{A}}$ may have huge coefficients, the *rays* of a linear projection of $\text{Amoeba}(\Delta_{\mathcal{A}})$ admit a concise description involving few bits, save for the transcendental coordinates coming from the “shifts” s_i . By applying our quantitative estimates from Section 2.3 we can then quickly find which chamber cone contains a given n -variate $(n+3)$ -nomial.

Theorem 3.7. *Following the notation of Theorem 3.4, suppose $f \in \mathcal{F}_{\mathcal{A}} \cap \mathbb{R}[x_1, \dots, x_n]$ and let τ denote the maximum bit size of any coordinate of \mathcal{A} . Then we can determine the unique chamber cone containing f — or correctly decide if f is contained in 2 or more chamber cones — within a number of arithmetic operations polynomial in $n + \tau$. Furthermore, if $f \in \mathcal{F}_{\mathcal{A}} \cap \mathbb{Z}[x_1, \dots, x_n]$, σ is the maximum bit size of any coefficient of f , and n is fixed, then we can also obtain a bit complexity bound polynomial in $\tau + \sigma$.* ■

Theorem 3.7 is the central tool behind our complexity results and follows immediately upon proving the correctness of (and giving suitable complexity bounds for) the following algorithm:

Algorithm 3.8.

Input: A subset $\mathcal{A} \subset \mathbb{Z}^n$ of cardinality $n+3$ (with \mathcal{A} not a pyramid and not contained in any $(n-1)$ -flat, and $\nabla_{\mathcal{A}}$ a hypersurface) and the coefficient vector c of an $f \in \mathcal{F}_{\mathcal{A}} \cap \mathbb{R}[x_1, \dots, x_n]$.

Output: Radiant subsets \mathcal{J} and \mathcal{J}' (corresponding to \mathcal{A}) generating the walls of the unique chamber cone containing f , or a true declaration that f is contained in at least 2 chamber cones.

Description:

- (-5) (Preprocessing) Compute the Hermite Factorization $H^T = U^T \hat{\mathcal{A}}$ and let B be the submatrix defined by the rightmost 2 columns of U .
- (-4) (Preprocessing) Let $\beta_1, \dots, \beta_{n+3}$ be the rows of B , $S := [s_{i,j}] := B \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} B^T$, and let s_i denote the row vector whose j^{th} coordinate is 0 or $\log |s_{i,j}|$ according as $s_{i,j}$ is 0 or not.
- (-3) (Preprocessing) Find all radiant subsets $\mathcal{J} \subset \{1, \dots, n+3\}$ corresponding to \mathcal{A} .
- (-2) (Preprocessing) For any radiant subset \mathcal{J} let $\beta'_{\mathcal{J}} := -\sum_{j \in \mathcal{J}} \beta_j$ and let $s_{\mathcal{J}}$ denote the row vector $s_i B$ for any fixed $i \in \mathcal{J}$.
- (-1) (Preprocessing) Sort the $\beta'_{\mathcal{J}}$ in order of increasing counter-clockwise angle with the x -coordinate ray and let R denote the resulting ordered collection of $\beta'_{\mathcal{J}}$.
- (0) (Preprocessing) For any radiant subset \mathcal{J} let $v_{\mathcal{J}} \in \mathbb{Q}^2$ denote the intersection of the lines $s_{\mathcal{J}} + \mathbb{R}\beta'_{\mathcal{J}}$ and $s_{\mathcal{J}'} + \mathbb{R}\beta'_{\mathcal{J}'}$, where $\beta'_{\mathcal{J}'}$ is the counter-clockwise neighbor of $\beta'_{\mathcal{J}}$.
- (1) Let $\text{ConeCount} := 0$.
- (2) Via binary search, attempt to find a pair of adjacent rays of the form $(v_{\mathcal{J}} + \mathbb{R}_+\beta'_{\mathcal{J}}, v_{\mathcal{J}} + \mathbb{R}_+\beta'_{\mathcal{J}'})$ containing $\text{Log}|c|B$.
 - (a) If ($\text{ConeCount} = 0$ and there is no such pair of rays)
or
($\text{ConeCount} = 1$ and there is such a pair of rays)
then output ‘‘Your f lies in at least 2 distinct chamber cones.’’ and STOP.
 - (b) If $\text{ConeCount} = 0$ and there is such a pair of rays, set $\text{ConeCount} := \text{ConeCount} + 1$, delete $\beta'_{\mathcal{J}}$ and $\beta'_{\mathcal{J}'}$ from R , and GOTO STEP (2).
- (3) Output ‘‘Your f lies in the unique chamber cone determined by \mathcal{J} and \mathcal{J}' .’’ and STOP.

Remark 3.9. An important detail for large scale computation is that the preprocessing steps (-5)–(0) need only be done once per support \mathcal{A} . This can significantly increase efficiency in applications where one has just one (or a few) \mathcal{A} and one needs to answer chamber cone membership queries for numerous f with the same support. ◇

Proof of Correctness of Algorithm 3.8: First note that the computed matrix B indeed has columns that form a basis for the right null-space of \mathcal{A} . This is because our assumptions on \mathcal{A} ensure that the rank of $\hat{\mathcal{A}}$ is $n+1$ and thus the last 2 rows of H^T consist solely of zeroes.

By construction, Theorem 3.4 then implies that the $\beta'_{\mathcal{J}}$ are exactly the reduced rays for $\nabla_{\mathcal{A}}$, modulo an invertible linear map. (The invertible map arises because right-multiplication by B induces an injective but non-orthogonal projection of the right null-space of $\hat{\mathcal{A}}$ onto \mathbb{R}^2 .)

It is then clear that the preprocessing steps do nothing more than give us a B suitable for Theorem 3.4 and a sorted set of reduced rays ready for chamber cone membership queries via binary search. In particular, since the reduced chamber cones cover \mathbb{R}^2 , the correctness of Steps (1)–(3) is clear and we are done. ■

In what follows, we will use the “soft-Oh” notation $O^*(f)$ to abbreviate bounds of the form $O(f(\log f)^{O(1)})$.

Complexity Analysis of Algorithm 3.8: Let us first approach our analysis from the more involved point of view of bit complexity. Our arithmetic complexity bound will then follow upon a quick revisit.

By Theorem 2.13 (which quotes bounds from [Sto00]), Step (-5) takes $O(n^{3.376}\tau^2)$ bit operations. Also, the resulting bit size for the entries of B is $O(n\tau)$.

The complexity of Step (-4) is negligible, save for the approximation of certain logarithms. The latter won’t come into play until we start checking on which side of a ray a point lies, so let us analyze the remaining preprocessing steps.

Step (-3) can be done easily through a greedy approach: one simply goes through the rows $\beta_2, \dots, \beta_{n+3}$ to see which rows are multiples of β_1 . Once this is done, one checks if the resulting set of indices is indeed radiant or not, and then one repeats this process with the remaining rows of B . In summary, this entails $O(n^2)$ arithmetic operations on number of bit-size $O(n\tau)$, yielding a total of $O^*(n^3\tau)$ bit operations.

Step (-2) has negligible complexity.

The comparisons in Step (-1) can be accomplished by computing the cosine and sine of the necessary angles via dot products and cross products. Via the well-known asymptotically optimal sorting algorithms, it is then clear that Step (-1) requires $O(n \log n)$ arithmetic operations on integers of bit size $O(n\tau)$, contributing a total of $O^*(n^2\tau \log n)$ bit operations.

Step (0) has negligible complexity.

At this point, we see that the complexity of the Preprocessing Steps (-5)–(0) is $O(n^{3.376}\tau^2)$ bit operations.

Continuing on to Steps (1)–(3), we now see that we are faced with $O(\log n)$ sidedness comparisons between a point and an oriented line. More precisely, we need to evaluate $O(\log n)$ signs of determinants of matrices of the form $\begin{bmatrix} \text{Log}|c|B - s_j \\ \beta'_j \end{bmatrix}$. Each such sign evaluation, thanks to Algorithm 2.15 and Lemma 2.16, takes $O(n30^{2n+5}L(\sigma + n\tau)L(\sigma)^{n+3}L(n\tau)^{n+2})$ bit operations.

We have thus proved our desired bit complexity bound which, while polynomial in $\tau + \sigma$ for fixed n , is visibly exponential in n . Note however that the exponential bottleneck occurs only in the sidedness comparisons of Step (2).

To obtain an improved arithmetic complexity bound, we then simply observe that the sidedness comparisons can be replaced by computations of signs of differences of monomials, simply by exponentiating the resulting linear forms in logarithms. Via recursive squaring [BS96, Thm. 5.4.1, pg. 103], it is then clear that each such comparison requires only $O(n^2\tau)$ arithmetic operations. So the overall number of arithmetic operations drops to polynomial in $n + \tau$ and we are done. ■

Let us now state some final combinatorial constructions before fully describing how chamber cones apply to real root counting.

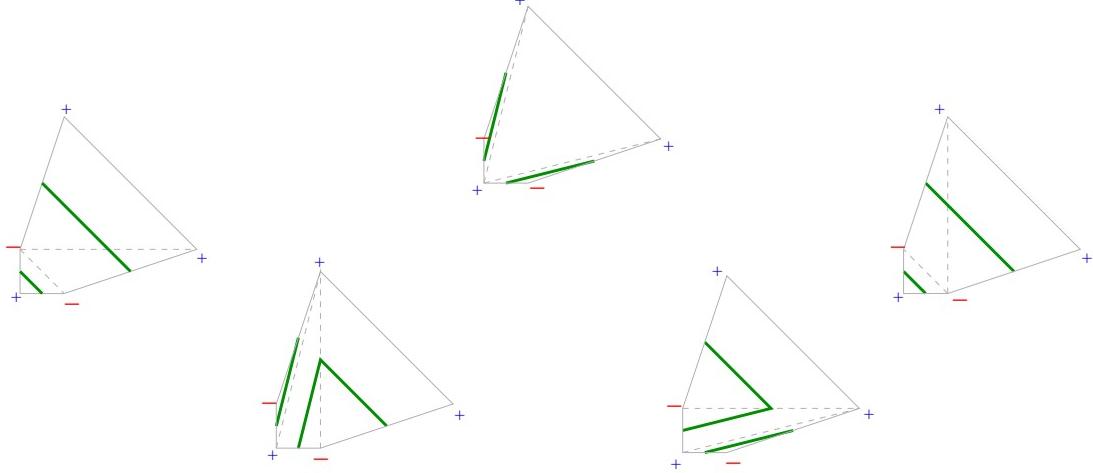
3.3. Canonical Viro Diagrams and the Probability of Lying in Outer Chambers. Our use of outer chambers and chamber cones enables us to augment an earlier construction of Viro. Let us first recall that a *triangulation* of a point set \mathcal{A} is simply a simplicial complex Σ whose vertices lie in \mathcal{A} .

Definition 3.10. We say that a triangulation of \mathcal{A} is coherent iff its maximal simplices are exactly the domains of linearity for some function ℓ that is convex, continuous, and piecewise linear on the convex hull of \mathcal{A} . In particular, we will sometimes define such an ℓ by fixing the values $\ell(a)$ for just those $a \in \mathcal{A}$ and then employing the convex hull of the points $(a, \ell(a))$ as ℓ ranges over \mathcal{A} . The resulting graph is known as the lower hull of the lifted point set $\{(a, \ell(a)) \mid a \in \mathcal{A}\}$. ◇

Definition 3.11. (See Proposition 5.2 and Theorem 5.6 of [GKZ94, Ch. 5, pp. 378–393].) Suppose $\mathcal{A} \subset \mathbb{Z}^n$ is finite and the convex hull of \mathcal{A} has positive volume and boundary $\partial\mathcal{A}$. Suppose also that \mathcal{A} is equipped with a coherent triangulation Σ and a function $s : \mathcal{A} \rightarrow \{\pm\}$ which we will call a distribution of signs for \mathcal{A} . We then call any edge with vertices of opposite sign an alternating edge, and we define a piece-wise linear manifold — the Viro diagram $\mathcal{V}_{\mathcal{A}}(\Sigma, s)$ — in the following local manner: For any n -cell $C \in \Sigma$, let L_C be the convex hull of the set of all midpoints of alternating edges of C , and then define $\mathcal{V}_{\mathcal{A}}(\Sigma, s) := \bigcup_{C \text{ an } n\text{-cell}} L_C \setminus \partial A$.

When $\mathcal{A} = \text{Supp}(f)$ and s is the corresponding sequence of coefficient signs, then we also call $\mathcal{V}_{\Sigma}(f) := \mathcal{V}_{\mathcal{A}}(\Sigma, s)$ the Viro diagram of f corresponding to Σ . ◇

Example 3.12. Consider $f(x) := 1 - x_1 - x_2 + 3x_1^4x_2 + 3x_1x_2^4$. Then $\text{Supp}(f) = \{(0, 0), (1, 0), (0, 1), (1, 4), (4, 1)\}$ and has convex hull a pentagon. So then there are exactly 5 coherent triangulations, yielding 5 possible Viro diagrams for f (drawn in thicker green lines):



Note that all these diagrams have exactly 2 connected components, with each component isotopic to an open interval. Note also that our f here is a 2-variate $(2+3)$ -nomial. ◇

Definition 3.13. Suppose $\mathcal{A} \subset \mathbb{Z}^n$ has cardinality $n+3$, is not a pyramid, is not contained in any $(n-1)$ -flat, and $\nabla_{\mathcal{A}}$ is a hypersurface. Also let B be any real $(n+3) \times 2$ matrix whose columns are a basis for the right null space of $\hat{\mathcal{A}}$. For any $f \in \mathcal{F}_{\mathcal{A}}$ let us then define $v(f) := (v_1(f), \dots, v_{n+3}(f)) := (\sum_{i \in \mathcal{I}} e_i) + (\sum_{j \in \mathcal{J}} e_j)$ where \mathcal{I} and \mathcal{J} are the unique radiant subsets corresponding to the unique chamber cone containing $\text{Log}|c|$. (We set $v(f) := \mathbf{0}$ should there not be a unique such chamber cone.) Let us then define $\text{ArchNewt}(f)$ to be the

convex hull of $\{(a_i, v_i) \mid i \in \{1, \dots, n+3\}\}$ and let Σ_f denote the triangulation of \mathcal{A} induced by the lower hull of $\text{ArchNewt}(f)$. Also, we call any polynomial of the form $\sum_{a_i \in Q} c_i x^{a_i} -$ with Q a cell of Σ_f — a canonical lower polynomial for f . Finally, we call $\mathcal{V}(f) := \mathcal{V}_\Sigma(f)$ the canonical Viro diagram of f . \diamond

Example 3.14. Let $f(x_1) := 1 - \frac{1}{2}x^{404} + x^{405} - 2x^{808}$, $c := (1, -\frac{1}{2}, 1, 2)$, and $\mathcal{A} := \{0, 404, 405, 808\}$. A routine calculation then reveals that $\{\{2\}, \{3\}\}$ is the pair of radiant subsets corresponding to the unique chamber cone containing $\text{Log}|c|$. We then obtain that $v(f) = (0, 1, 1, 0)$ and thus $\text{ArchNewt}(f)$ is  (modulo some artistic stretching). In particular, Σ_f thus has the single cell $[0, 808]$, which is an alternating cell, and thus $\mathcal{V}(f)$ consists of a single point. More than coincidentally, f has exactly 1 positive root. \diamond

Example 3.15. Returning to Example 3.12, let $c := (1, -1, -1, 3, 3)$. A routine calculation then reveals that the unique chamber cone containing $\text{Log}|c|$ is defined by the pair of radiant subsets $\{\{2\}, \{3\}\}$. So then $v(f) = (0, 1, 1, 0, 0)$ and Σ_f is the exactly the upper middle triangulation from the illustration of Example 3.12. $\mathcal{V}(f)$ then consists of 2 disjoint open intervals and, more than coincidentally, the positive zero set of f has exactly 2 connected components, each homeomorphic to an open interval. \diamond

Theorem 3.16. Following the notation above, set $\hat{f}_t(x) := \sum_{i=1}^{n+3} c_i t^{v_i(f)} x^{a_i}$ and assume in addition that $\mathcal{A} \cap Q$ has cardinality n for all facets Q of $\text{Conv}\mathcal{A}$. Then $c := (c_1, \dots, c_{n+3})$ lies in an outer chamber \implies the positive zero sets of \hat{f}_t , as t ranges over $(0, 1]$, are each diffeotopic to the positive zero set of \hat{f}_1 . In particular, $\hat{f}_1 = f$ and thus, when c lies in an outer chamber, the positive zero set of f is isotopic to $\mathcal{V}(f)$.

Remark 3.17. For $n=1$ we thus obtain that the number of positive roots of the tetranomial f is exactly the cardinality of its canonical Viro diagram. \diamond

Proof: By construction, the image of $\text{Log}((c_1 t^{v_1(f)}, \dots, c_{n+3} t^{v_{n+3}(f)}))$ as t ranges over $(0, 1]$ is simply a ray entirely contained in a unique chamber cone. Moreover, by assumption (and since outer chambers are log convex), the ray is also contained entirely in $\text{Log}|\cdot|$ of an outer chamber. So the first part of our theorem follows from Lemma 2.10.

The final part of our theorem is then just a reformulation of Viro's Theorem on the isotopy type of toric deformations of real algebraic sets (see, e.g., [GKZ94, Thm. 5.6]). \blacksquare

The main contribution of our paper is thus an efficient method to associate a *canonical* Viro diagram to the positive zero set of a given f so that both C^1 manifolds have the same topology. Such a method appears to be new, although the necessary ingredients have existed in the literature since at least the 1990s. In particular, to the best of our knowledge, all earlier applications of Viro's method designed clever f having the same topology as some specially tailored Viro diagram, thus going in the opposite direction of our construction.

We state up front that our method does *not* work for all f . However, our development yields a sufficient condition — outer chamber membership — that holds with high probability under the stable log-uniform measure.

Theorem 3.18. Suppose $\mathcal{A} \subset \mathbb{Z}^n$ has cardinality $n+3$, is not a pyramid, is not contained in any $(n-1)$ -flat, and $\nabla_{\mathcal{A}}$ is a hypersurface. Suppose also that the coefficients of an $f \in \mathcal{F}_{\mathcal{A}} \cap \mathbb{R}[x_1, \dots, x_n]$ are independently chosen via the stable log-uniform measure over \mathbb{R} (resp. \mathbb{Z}). Then with probability 1, f lies in some outer chamber. In particular, if we assume

in addition that $\mathcal{A} \cap Q$ has cardinality n for all facets Q of $\text{Conv}\mathcal{A}$, then the positive zero set of f is isotopic to $\mathcal{V}(f)$ with probability 1.

Proof: By Theorem 3.4, $\text{Amoeba}(\Delta_{\mathcal{A}})$ is an n -plane bundle over $\text{Amoeba}(\overline{\Delta}_{\mathcal{A}})$, where $\overline{\Delta}_{\mathcal{A}} \in \mathbb{Z}[a, b]$ and $\Delta_{\mathcal{A}}(c_1, \dots, c_{n+3}) = \overline{\Delta}_{\mathcal{A}}(\alpha(c), \beta(c))$ for suitable monomials α and β in c . Furthermore, thanks to Corollary 8 of [PST05], the complement of $\text{Amoeba}(\overline{\Delta}_{\mathcal{A}})$ has no bounded convex connected components.

Letting c denote the coefficient vector of f we thus obtain that f lies in an outer chamber iff $\text{Log}|c| \notin \text{Amoeba}(\Delta_{\mathcal{A}})$. In particular, it is clear by the Passare-Rullgård Theorem that in any large centered cube C , the volume of $\text{Amoeba}(\Delta_{\mathcal{A}}) \cap C$ occupies a vanishingly small fraction of C . So the first assertion is proved.

The final assertion is an immediate consequence of the first assertion and Theorem 3.16. ■

Theorem 1.4 then follows easily from Theorems 3.16 and 3.18. The applications of Theorems 3.16 and 3.18 to computational real topology will be pursued in another paper.

3.4. Proving Theorem 1.4. Let us first consider the following algorithm for counting the positive roots of “most” real univariate tetranomial.

Algorithm 3.19.

Input: A tetranomial $f \in \mathbb{R}[x_1]$ with support \mathcal{A} .

Output: A number in $\{0, 1, 2, 3\}$ that, for any f in an outer chamber of $\nabla_{\mathcal{A}}$, is exactly the number of positive roots of f .

Description:

- (1) Via Algorithm 3.8, and any sub-quadratic planar convex hull algorithm (see, e.g., [OSvK00]), compute the canonical Viro diagram $\mathcal{V}(f)$.
- (2) If f did not lie in a unique chamber cone then output ‘‘Your f does not lie in an outer chamber so please use an alternative method.’’ and STOP.
- (3) Output the cardinality of $\mathcal{V}(f)$ and STOP.

Assuming Algorithm 3.19 is correct, we can count the real roots of f simply by applying Algorithm 3.19 to $f(x_1)$ and $f(-x_1)$. Theorem 1.4 thus follows immediately upon proving the correctness of our last algorithm and a suitable complexity bound. We now do this.

Proof of Correctness of Algorithm 3.19: By Theorem 3.16, the number of positive roots of f is exactly the cardinality of $\mathcal{V}(f)$ whenever f is in an outer chamber. So we indeed have correctness. ■

Complexity Analysis of Algorithm 3.19: Let us first observe that Algorithm 3.19 indeed gives a correct answer with probability 1 (relative to the stable log-uniform measure): this is immediate from Theorem 3.18.

So now we need only prove a suitable complexity bound. Let us begin with the more refined setting of bit complexity: From our earlier complexity analysis of Algorithm 3.8, it is clear that Step (1) needs $O(\log^2 D) + O(L(h + \log D)L(h)^4 L(\log D)^3)$ bit operations, neglecting the computation of $\mathcal{V}(f)$. The complexity of computing $\mathcal{V}(f)$ (which is essentially dominated by computing the convex hull of 4 points with coordinates of bit size $O(\log D)$) is clearly negligible in comparison. So we obtain a final bit complexity bound of $O^*((h + \log D)h^4 \log^3 D)$. (The complexity of Steps (2) and (3) is negligible.)

For arithmetic complexity, our earlier complexity analysis of Algorithm 3.8 specializes easily to an upper bound of $O(\log^2 D)$. (The speed-up arises from the easiness of checking

inequalities involving integral powers of real numbers in the BSS model over \mathbb{R} .) So we are done. ■

Remark 3.20. *It is important to note that when f lies in a chamber cone but not in any outer chamber, Algorithm 3.19 can give a wrong answer. However, thanks to Theorem 3.18, such an occurrence has probability 0 under the stable log-uniform measure. ◇*

4. PROVING THEOREM 1.5

To prepare for the proof, we first set up some notation. Fix positive integers ℓ and m . Let $P = (p_{ij}) \in \mathbb{N}^{\ell \times m}$ be a matrix of natural numbers ordered as

$$p_{11} \geq p_{21} \geq \cdots \geq p_{s1} \quad \text{and} \quad p_{ij} > p_{i(j+1)},$$

for $i = 1, \dots, s$ and $j = 1, \dots, m - 1$. Also, let a_{ij} be indeterminates over the same indexing set. Consider now the following polynomial:

$$\begin{aligned} S_P(x) &:= \sum_{i=1}^{\ell} \left(\sum_{j=1}^m a_{ij} x^{p_{ij}} \right)^2 \in \mathbb{N}[a_{ij}][x] \\ &= g_d(P)x^d + \cdots + g_1(P)x + g_0(P), \end{aligned} \tag{4.1}$$

in which each nonzero $g_i(P)$ is a homogeneous (quadratic) polynomial in $\mathbb{N}[a_{ij}]$. Notice that at most ℓm^2 of these g_i are nonzero. We will refer to the integer p_{ij} as the *exponent* corresponding to the *coefficient* a_{ij} .

Lemma 4.1. *The set of polynomials $g_i(P)$ arising from $S_P(x)$ as P ranges over all $\ell \times m$ nonnegative integer matrices is finite.*

Proof: The form of Equation (4.1) allows for at most $2^{\ell m^2}$ possible such g_i . ■

Suppose now that $f \in \mathbb{R}[x]$ has degree d and is a sum of ℓ squares, each involving at most m terms. Then, there is a set of exponents P and an assignment $\bar{a}_{ij} \in \mathbb{R}$ for the coefficients a_{ij} such that $f = S_P(x)$. If we fix a set of exponents P for f , then conversely, any real point in the variety determined by the g_i gives a representation of f as a sum of ℓ squares, each involving at most m terms.

We will prove Theorem 1.5 using contradiction by showing that a certain infinite family of trinomials cannot all have sparse representations of the form in Equality (4.1). For this approach to work, however, we will need to find a universal set of coefficients \bar{a}_{ij} that “represents” all of them.

Lemma 4.2. *Let $F \subset \mathbb{R}[x]$ be an infinite collection of polynomials which are sums of ℓ squares, each involving at most m terms. Moreover, suppose that the nonzero coefficients of polynomials $f \in F$ come from a finite set C . Then, there is an infinite subset $f_1, f_2, \dots \in F$ of them with corresponding exponents $P_1, P_2, \dots \in \mathbb{N}^{\ell \times m}$ and a single set of real coefficients \bar{a}_{ij} such that $f_k = S_{P_k}(x)|_{a_{ij}=\bar{a}_{ij}}$ for all k .*

Proof: Let $f_1, f_2, \dots \in F$ be an infinite sequence of polynomials with corresponding exponents $P_1, P_2, \dots \in \mathbb{N}^{\ell \times m}$. Also let S be the set of all possible (polynomial) coefficients of polynomials $S_{P_k}(x)$; from Lemma 4.1, this set is finite. By assumption, each f_k gives rise to a set of equations involving a subset of S and real numbers coming from the finite set C . The set of all such equations is again finite, and therefore, by the infinite pigeon-hole

principle, there is a subsequence of the f_k which have the same set of equations governing their coefficients a_{ij} . Picking any real solution to such a set of equations finishes the proof.

We now have all the tools we need to prove Theorem 1.5. However, to make the argument less cumbersome, we first recall “little-oh” notation: for a function $h : \mathbb{N} \rightarrow \mathbb{R}$, we say that $h(n) = o(n)$ if

$$\lim_{n \rightarrow \infty} \frac{h(n)}{n} = 0.$$

It is easy to see that the sum of any finite number of such functions is also $o(n)$. Moreover, if $\lim_{n \rightarrow \infty} \frac{p(n)}{n} = p$ for some constant p , then $p(n) = np + o(n)$.

Proof of Theorem 1.5: Suppose, by way of contradiction, that every positive definite trinomial can be written sparsely as the sum of ℓ squares, each involving at most m terms. Consider the infinite sequence of positive definite trinomials,

$$f_k = x^{2k} + x^{2k-1} + 1, \quad k = 1, 2, \dots \quad (4.2)$$

Using Lemma 4.2, we can find a subsequence f_{k_s} with corresponding exponents $P_{k_1}, P_{k_2}, \dots \in \mathbb{N}^{\ell \times m}$ and a single set of real numbers \bar{a}_{ij} such that $f_{k_s} = S_{P_{k_s}}(x)|_{a_{ij}=\bar{a}_{ij}}$ for $s = 1, 2, \dots$. We will assume that the \bar{a}_{ij} are chosen such that the number of them that are zero is maximal among all such sets of coefficients \bar{a}_{ij} . Moreover, we shall assume that this maximality still holds if we restrict to any infinite subsequence of the f_{k_s} . For clarity of exposition in what follows, we will not keep updating the subscripting of indices when taking subsequences.

Given an exponent matrix $P_{k_s} \in \mathbb{N}^{\ell \times m}$, define a new matrix

$$\tilde{P}_{k_s} = \frac{1}{k_s} P_{k_s}.$$

This corresponds naturally to the transformation $x \mapsto x^{1/k_s}$ applied to both sides of an equation $f_{k_s}(x) = S_{P_{k_s}}(x)$. Since $\deg(f_{k_s}) = 2k_s$, each matrix \tilde{P}_{k_s} has entries in the interval $[0, 1]$. By compactness, we may choose a subsequence P_{k_s} such that \tilde{P}_{k_s} converges in the (entry-wise) Euclidean norm to a matrix $\tilde{P} = (\tilde{p}_{ij}) \in [0, 1]^{\ell \times m}$. Clearly, we have $\tilde{p}_{11} = 1$ and also that some entry of \tilde{P} is 0; it turns out that these are the only possible numbers for entries of \tilde{P} which play a role in (4.1).

Claim: If $\tilde{p}_{ij} \neq 0, 1$, then $\bar{a}_{ij} = 0$.

Suppose, on the contrary, that \tilde{P} contains $r > 2$ real numbers, p_1, \dots, p_r with $0 = p_1 < \dots < p_r = 1$. Each power of x occurring in a summand of Equality (4.1), after squaring, is of the form

$$k_s p_u + k_s p_v + o(k_s). \quad (4.3)$$

Thus, for all sufficiently large s , the powers of x occurring in Equality (4.1) are partitioned into classes determined by the number of distinct values,

$$p_u + p_v, \quad u, v = 1, \dots, r.$$

Notice that the numbers in Formula (4.3) all become strictly smaller (resp. larger) than $2k_s - 1$ (resp. 0) as $s \rightarrow \infty$ unless $u = v = r$ (resp. $u = v = 1$). In particular, with s large, the polynomials $g_w(P_{k_s}) \in \mathbb{N}[a_{ij}]$ for

$$w = 2k_s + o(k_s) \quad \text{and} \quad w = 0 + o(k_s) \quad (4.4)$$

do not involve the indeterminates a_{ij} coming from exponents of the form $k_s p_u + o(k_s)$ with $u \neq 1, r$. Conversely, each monomial in every polynomial $g_w(P_{k_s})$ with w not in one of the classes from (4.4) contains at least one such indeterminate a_{ij} .

Since the only nonzero coefficients of the sequence (4.2) come from the classes from (4.4), it follows that we may replace with 0 all coefficients \bar{a}_{ij} corresponding to exponents $k_s p_u + o(k_s)$ with $u \neq 1, r$ and still have an equality

$$f_{k_s} = S_{P_{k_s}}(x)|_{a_{ij}=\bar{a}_{ij}}.$$

The claim therefore follows from the maximality property of the chosen set of coefficients \bar{a}_{ij} .

We next examine what all this says about the limiting behavior of the expressions from Equality (4.1). From the claim, it follows that when s is large, we need only consider those exponents from the matrices P_{k_s} that are on the order

$$k_s + o(k_s) \text{ and } 0 + o(k_s).$$

Fix such a large s and consider the smallest exponent p on the order $k_s + o(k_s)$ that occurs with a nonzero coefficient in (4.1) after substituting the \bar{a}_{ij} for the a_{ij} . When the sum from (4.1) is expanded, the term x^{2p} will appear with positive coefficient; i.e.,

$$g_{2p}(P_{k_s})|_{a_{ij}=\bar{a}_{ij}} \neq 0.$$

It follows from the form of the sequence (4.2) that $p = k_s$ (since the other nonzero term is odd). In particular, it is not possible to obtain a nonzero coefficient for x^{2k_s-1} in f_{k_s} . This contradiction completes the proof. ■

REFERENCES

- [AI11] Avendaño, Martín and Ibrahim, Ashraf, “*Multivariate Ultrametric Root Counting*,” preprint, Texas A&M University, submitted for publication.
- [AAR11] Ascher, Kenneth; Avendaño, Martín; and Rojas, J. Maurice, “*Smale’s 17th Problem over the Reals in One Variable*,” in progress.
- [BS96] Bach, Eric and Shallit, Jeff, *Algorithmic Number Theory, Vol. I: Efficient Algorithms*, MIT Press, Cambridge, MA, 1996.
- [Bak77] Baker, Alan, “*The Theory of Linear Forms in Logarithms*,” in Transcendence Theory: Advances and Applications: proceedings of a conference held at the University of Cambridge, Cambridge, Jan.–Feb., 1976, Academic Press, London, 1977.
- [BPR06] Basu, Saugata; Pollack, Ricky; and Roy, Marie-Francoise, *Algorithms in Real Algebraic Geometry*, Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, 2006.
- [Ber03] Bernstein, Daniel J., “*Computing logarithm intervals with the arithmetic-geometric mean iterations*,” available from <http://cr.ypr.to/papers.html>, 2003.
- [BRS09] Bihan, Frederic; Rojas, J. Maurice; and Stella, Casey E., “*Faster Real Feasibility via Circuit Discriminants*,” proceedings of ISSAC 2009 (July 28–31, Seoul, Korea), pp. 39–46, ACM Press, 2009.
- [BSZ00] Bleher, Pavel; Shiffman, Bernard; and Zelditch, Steve, “*Universality and scaling of correlations between zeros on complex manifolds*,” Invent. Math. 142 (2000), no. 2, pp. 351–395.
- [BCSS98] Blum, Lenore; Cucker, Felipe; Shub, Mike; and Smale, Steve, *Complexity and Real Computation*, Springer-Verlag, 1998.
- [DFS07] Dickenstein, Alicia; Feichtner, Eva Maria; and Sturmfels, Bernd, “*Tropical Discriminants*,” J. Amer. Math. Soc., **20** (2007), pp. 1111–1133.
- [DRRS07] Dickenstein, Alicia; Rojas, J. Maurice; Rusek, Korben; Shih, Justin, “*Extremal Real Algebraic Geometry and A-Discriminants*,” Moscow Mathematical Journal, vol. 7, no. 3, (July–September, 2007), pp. 425–452.
- [GKZ94] Gel’fand, Israel Moseyevitch; Kapranov, Misha M.; and Zelevinsky, Andrei V.; *Discriminants, Resultants and Multidimensional Determinants*, Birkhäuser, Boston, 1994.

- [Hab48] Habicht, Walter, “*Eine Verallgemeinerung des Sturmschen Wurzelzählverfahrens*,” Comment. Math. Helv. **21** (1948), pp. 99–116.
- [HTZEKM09] Hemmer, Michael; Tsigaridas, Elias P.; Zafeirakopolous, Zafeirakis; Emiris, Ioannis Z.; Karavelas, Menelaos I.; and Mourrain, Bernard, “*Experimental evaluation and cross-benchmarking of univariate real solvers*,” proceedings of SNC 2009 (Symbolic-Numeric Computation, Kyoto, Japan, August 2–5), pp. 45–54, ACM Press, 2009.
- [Kap91] Kapranov, Misha, “*A characterization of A-discriminantal hypersurfaces in terms of the logarithmic Gauss map*,” Mathematische Annalen, 290, 1991, pp. 277–285.
- [KM09] Kojima, M. and Muramatsu, M., “*A Note on Sparse SOS and SDP Relaxations for Polynomial Optimization Problems over Symmetric Cones*,” Computational Optimization and Applications Vol. 42 (1), pp. 31–41 (2009).
- [Kos88] Kostlan, Eric J., “*Complexity theory of numerical linear algebra*,” Journal of Computational and Applied Mathematics Volume 22, Issues 2-3, June 1988, pp. 219–230
- [Las06] Lassere, Jean B., “*Convergent SDP-Relaxations in Polynomial Optimization with Sparsity*,” SIAM J. Optim., Vol. 17, No. 3, pp. 822–843.
- [Las09] Lasserre, Jean-Michel, “*Moments and sums of squares for polynomial optimization and related problems*,” Journal of Global Optimization, vol. 45, no. 1, pp. 39–61, sept. 2009.
- [LLL82] Lenstra, Arjen K.; Lenstra (Jr.), Hendrik W.; and Lovász, László, “*Factoring Polynomials with Rational Coefficients*,” Math. Ann. 261 (1982), no. 4, pp. 515–534.
- [LM01] Lickteig, Thomas and Roy, Marie-Francoise, “*Sylvester-Habicht Sequences and Fast Cauchy Index Computation*,” J. Symbolic Computation (2001) **31**, pp. 315–341.
- [MR04] Malajovich, Gregorio and Rojas, J. Maurice, “*High Probability Analysis of the Condition Number of Sparse Polynomial Systems*,” Theoretical Computer Science, special issue on algebraic and numerical algorithms, Vol. 315, no. 2–3, (May 2004), pp. 525–555.
- [Nes03] Nesterenko, Yuri, “*Linear forms in logarithms of rational numbers*,” Diophantine approximation (Cetraro, 2000), pp. 53–106, Lecture Notes in Math., 1819, Springer, Berlin, 2003.
- [OSvK00] Overmars, Mark; Schwarzkopf, Otfried; and van Kreveld, Marc, *Computational Geometry: Algorithms and Applications*, Springer Verlag, 2000.
- [Par03] Parrilo, Pablo A., “*Semidefinite programming relaxations for semialgebraic problems*,” Algebraic and geometric methods in discrete optimization, Math. Program. 96 (2003), no. 2, Ser. B, pp. 293–320.
- [PR04] Passare, Mikael and Rullgård, Hans, “*Amoebas, Monge-Ampère measures, and triangulations of the Newton polytope*,” Duke Math. J. Vol. 121, No. 3 (2004), pp. 481–507.
- [PT05] Passare, Mikael and Tsikh, August, “*Amoebas: their spines and their contours*,” Idempotent mathematics and mathematical physics, Contemp. Math., v. 377, Amer. Math. Soc., Providence, RI, 2005, pp. 275–288.
- [PST05] Passare, Mikael; Sadykov, Timur; and Tsikh, August, “*Singularities of hypergeometric functions in several variables*,” Compositio Math. 141 (2005), pp. 787–810.
- [PRT09] Pébay, Philippe; Rojas, J. Maurice; Thompson, David C., “*Optimization and NP_R-Completeness of Certain Fewnomials*,” proceedings of SNC 2009 (August 3–5, 2009, Kyoto, Japan), pp. 133–142, ACM Press, 2009.
- [PRRT11] Pébay, Philippe; Rojas, J. Maurice; Rusek, Korben; and Thompson, David C., “*Simple Homotopies for Just the Real Roots of Polynomial Systems*,” preprint, Sandia National Laboratories, 2011.
- [Pou71] Pourchet, Y., “*Sur la représentation en somme de carrés des polynômes à une indéterminée sur un corps de nombres algébriques*,” Acta Arithm. 19 (1971), pp. 89–104.
- [RS02] Rahman, Q. I. and Schmeisser, G., *Analytic Theory of Polynomials*, London Mathematical Society Monographs 26, Oxford Science Publications, 2002.
- [RY05] Rojas, J. Maurice and Ye, Yinyu, “*On Solving Sparse Polynomials in Logarithmic Time*,” Journal of Complexity, special issue for the 2002 Foundations of Computation Mathematics (FOCM) meeting, February 2005, pp. 87–110.
- [SS96] Shub, Mike and Smale, Steve, “*The Complexity of Bezout’s Theorem IV: Probability of Success; Extensions*,” SIAM J. Numer. Anal., **33** (1996), no. 1, pp. 128–148.
- [Sto00] Storjohann, Arne, “*Algorithms for Matrix Canonical Forms*,” doctoral dissertation, Swiss Federal Institute of Technology, Zurich, 2000.

[Stu35] Sturm, Jacques Charles-François, “*Mémoire sur la résolution des équations numériques*,” Inst. France Sc. Math. Phys., **6** (1835).

HARVARD UNIVERSITY, MASSACHUSETTS HALL, CAMBRIDGE, MA 02138
E-mail address: hypo3400@gmail.com

REDWOOD CENTER FOR THEORETICAL NEUROSCIENCE, 575A EVANS HALL, MC 3198 BERKELEY, CA 94720-3198.

E-mail address: chillar@msri.org

MIT, 77 MASS. AVE., CAMBRIDGE, MA 02139
E-mail address: dpopov@mit.edu

TAMU 3368, TEXAS A&M UNIVERSITY, COLLEGE STATION, TX 77843-3368
E-mail address: rojas@math.tamu.edu